

Multi-Petabyte Data Distribution in Industry & Science with CernVM File System

Georgios Christodoulis, CERN
Andriy Utkin




CernVM
File system

Software distribution in Particle Physics



European Environment for Scientific Software Installations (EESSI)



CernVM
File system


CernVM Filesystem (CVMFS)

```
[root@christge-cl ~]# tree -L 2 /cvmfs/atlas.cern.ch/
/cvmfs/atlas.cern.ch/
└── repo
    ├── ATLASLocalRootBase
    ├── benchmarks
    ├── conditions -> /cvmfs/atlas-condb.cern.ch/repo/conditions
    ├── containers
    ├── dev
    ├── sw
    ├── test
    ├── tools
    └── tutorials
└── test1
```


10 directories, 2 files

```
[root@christge-cl ~]# cat /cvmfs/atlas.cern.ch/test1 #JIT download test1
test
[root@christge-cl ~]#
```

CVMFS is a filesystem in user space




CVMFS is a filesystem in user space



File distribution


- Files are stored by content using hash (Merkle) trees
- Allows for data verification and compression
- LRU cache on client side
- Content addressable storage allows **de-duplication** (each hash is stored only once)



CVMFS at CERN with numbers


- > 33B files
- > 2PB of data
- > 8k container images
- ~ 290 Repositories

A well established tool for
software distribution
(and preservation)!



Distribution of Container Images


Images are **unseekable**
and **unindexed** file
archives.



[K. Tokunaga]



Lazy pulling

- pull/download the essential parts of the image on demand.
- The container can start almost immediately (possibly slower execution)



All fallback to legacy pulling if image not available in required format.

unpacked.cern.ch



>12 TB of compressed and deduplicated images
scattered among >120million files.

(~1.3PB logical data)